

Fast Markov chain Monte Carlo sampling for sparse Bayesian inference in high-dimensional inverse problems using L1-type priors

Felix Lucka^{1,2}

¹ Institute for Computational and Applied Mathematics, University of Münster, Einsteinstr. 62, D-48149 Münster, Germany

² Institute for Biomagnetism and Biosignalanalysis, University of Münster, Malmedyweg 15, D-48149 Münster, Germany

E-mail: felix.lucka@uni-muenster.de

Abstract. Sparsity has become a key concept for solving of high-dimensional inverse problems using variational regularization techniques. Recently, using similar sparsity-constraints in the Bayesian framework for inverse problems by encoding them in the prior distribution has attracted attention. Important questions about the relation between regularization theory and Bayesian inference still need to be addressed when using sparsity promoting inversion. A practical obstacle for these examinations is the lack of fast posterior sampling algorithms for sparse, high-dimensional Bayesian inversion: Accessing the full range of Bayesian inference methods requires being able to draw samples from the posterior probability distribution in a fast and efficient way. This is usually done using Markov chain Monte Carlo (MCMC) sampling algorithms. In this article, we develop and examine a new implementation of a single component Gibbs MCMC sampler for sparse priors relying on L1-norms. We demonstrate that the efficiency of our Gibbs sampler increases when the level of sparsity or the dimension of the unknowns is increased. This property is contrary to the properties of the most commonly applied Metropolis-Hastings (MH) sampling schemes: We demonstrate that the efficiency of MH schemes for L1-type priors dramatically decreases when the level of sparsity or the dimension of the unknowns is increased. Practically, Bayesian inversion for L1-type priors using MH samplers is not feasible at all. As this is commonly believed to be an intrinsic feature of MCMC sampling, the performance of our Gibbs sampler also challenges common beliefs about the applicability of sample based Bayesian inference.

AMS classification scheme numbers: 65J22, 62F15, 65C05, 65C60

Submitted to: *Inverse Problems*

1. Introduction

1.1. Sparse Bayesian inversion

Solving high-dimensional inverse problems using *sparsity* constraints as *a priori information* has led to enormous advances in various application areas. *Total variation* (TV) deblurring [47, 6] uses sparsity constraints on the gradient of the unknown quantity and is successfully used in many imaging applications. By the notion of *compressed sensing* [14, 15], a number of techniques are summarized, which rely on the idea that high quality reconstructions can be obtained from a small amount of data, if a sparse basis for the unknowns is a priori known. Traditionally, sparsity constraints are formulated in the framework of *variational regularization* when introduced as a priori information to inverse problems. One popular approach is to use regularization functionals incorporating L1 norms. However, this type of sparsity constraints can also be formulated and examined in the framework of Bayesian statistics [33]. Addressing high dimensional, ill-posed inverse problems as problems of Bayesian inference has gained growing attention over the years [50, 20, 31, 30]. It allows an easy formulation of *a priori information* on the solution via *a priori probability distributions (prior)*. Furthermore, a specific inference strategy that is called the *maximum a posteriori* estimate (MAP) corresponds to variational regularization (the prior corresponds to the regularization functional). In general, the Bayesian solution to an inverse problem is given by the *a posteriori probability distribution (posterior)* over the parameter space (the MAP estimate is the point maximizing this distribution). The analysis of sparse Bayesian inversion is far less elaborate up to now, and a number of exciting questions still remain to be addressed. In particular, sparse inversion is an interesting topic to study the relation between regularization theory and Bayesian inference. This relation is well understood for regularization using L2 norms, which corresponds to Bayesian inference with Gaussian priors, see, e.g., [31]. While the differences in these scenarios are subtle, they become way more pronounced in the context of sparse inversion using *L1-type priors*, i.e., priors, which rely on L1 norms of the unknowns [36, 35]. A central tool to study these differences is the examination of the posterior by *Monte Carlo sampling* methods. The standard sampling techniques were designed for Bayesian inference in low-dimensional, well-posed problems and often fail when used in scenarios arising from typical inverse problems. For these reasons, a number of specific sampling techniques for ill-posed, high-dimensional problems have already been developed [28, 22, 23, 41, 44]. However, these techniques mainly address Gaussian priors. For the sparsity-promoting L1-type priors they may fail dramatically. This observation is in line with the fact that efficient optimization techniques for the corresponding variational regularization schemes are still a vital field of research as well [21, 5].

1.2. Contributions and Structure

Our work was motivated by questions that arise when Bayesian inference using general L1-type priors is applied to edge preserving image reconstruction, similar to the scenarios discussed in [36, 35, 33]. A major problem we and others faced was that the conventional *Markov chain Monte Carlo (MCMC)* tools for sampling the posterior distribution fail in such situations, rendering many examinations infeasible. Therefore, we develop and examine new and more efficient implementations of sampling algorithms for these situations first. In this paper, we present a fast

implementation of a Gibbs sampling algorithm and study its performance in two typical inverse problems scenarios. Thereby, we provide a solid basis for addressing more sophisticated questions in sparse Bayesian inversion in the future.

In Section 2, we describe the setting and methods used. Detailed numerical examinations of all MCMC algorithms for two test scenarios are presented in Section 3. In Section 4, the results are discussed and we point to future directions of development. Additionally, implementation details and code for the new sampling algorithms is provided in Appendix A.

2. Methods

In this section, we will first introduce the general setting for our examinations (Section 2.1) and the basics of Bayesian inference. Then we review the basic principles of MCMC-based posterior inference and present the most popular MCMC sampling schemes, the *Metropolis-Hastings* algorithm and the *Gibbs sampling* algorithm (Section 2.2). The intention of these first two sections is to make the article more accessible for readers which mainly used variational regularization techniques so far and have little experience with Bayesian techniques. The more advanced reader can skip these sections. Section 2.3 contains the main contributions of this article, i.e., the development of a *Gibbs sampling* scheme for L1-type priors which relies on a robust numerical implementation of an exact, explicit sampling from the conditional single component posterior. In the last section (Section 2.4), we will explain the methods used for the evaluation of the sampling performance in the computational studies. The experienced reader may, again, skip this section.

2.1. General Setting and Bayesian Formulation

In general, we consider the inverse problem of solving a continuous, linear, ill-posed operator equation. Here, we start from the following discrete model chosen for obtaining a computational solution (the *computational model*):

$$m = Au + \varepsilon, \quad (1)$$

where $m \in \mathbb{R}^k$ represents the given measurement data, $u \in \mathbb{R}^n$ represents the unknowns derived from a discretization of the computational domain, $A \in \mathbb{R}^{k \times n}$ is the discretization of the continuous forward operator with respect to the domains of u and m and $\varepsilon \in \mathbb{R}^k$ is an additive, stochastic noise term. Accounting for the stochastic nature of the noise term renders (1) into a relation between the k -dim random variables M and \mathcal{E} (the *likelihood model*):

$$M = Au + \mathcal{E} \quad (2)$$

See [29, 3] for details on the implications of this step. For simplicity, we assume $\mathcal{E} \sim \mathcal{N}(0, \sigma^2 I_k)$, $\sigma > 0$ here, where I_k is the k -dim identity matrix (the extension to general Gaussian noise is straight forward). Now, the conditional probability density of M given u is determined by (2) and is, thus, called the *likelihood density*:

$$p_i(m|u) = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{k}{2}} \exp \left(-\frac{1}{2\sigma^2} \|m - Au\|_2^2 \right) \quad (3)$$

Due to the ill-posedness of (1), inference about u given M on the basis of (3) is not feasible with standard statistical inference strategies. *Bayesian inference strategies*

rely on considering u as a random variable itself (U in our notation) and on encoding a priori information about U in its density, $p_{pr}(u)$, which is therefore called the *prior*. Then, the model can be inverted using Bayes' rule:

$$p_{post}(u|m) = \frac{p_{li}(m|u)p_{pr}(u)}{p(m)} \quad (4)$$

The conditional density of U given M is called the *posterior*. In Bayesian inference, this density is the complete solution to the inverse problem. The term $p(m)$ is called the *model-evidence* and for our aims, it is just a normalizing constant, which is of no further importance. There are several ways to exploit the information about U contained in the posterior. The most popular one, called the *maximum a posteriori* estimate (*MAP*), is to infer a point estimate for U by searching for the highest *mode* of the posterior. Another way to obtain a point estimate, called the *conditional mean* estimate (*CM*), is to compute the mean/expected value of the posterior:

$$\hat{u}_{\text{MAP}} := \operatorname{argmax}_{u \in \mathbb{R}^n} \{ p_{post}(u|m) \} \quad (5)$$

$$\hat{u}_{\text{CM}} := \mathbb{E}[u|m] = \int u p_{post}(u|m) du \quad (6)$$

Practically, computing the MAP estimate is a high-dimensional *optimization* problem, whereas computing the CM estimate is a high-dimensional *integration* problem. Apart from point estimates, computing *confidence intervals*, *conditional covariance* or *histogram* estimates are other applications of posterior-based inference. See [31] for an overview and, e.g., [22, 23, 7] for the applications to remote sensing, algae population dynamics and image deblurring.

This far, we did not specify the concrete form of the prior $p_{pr}(u)$, which is actually the most important step within the Bayesian formalism. A common choice linking Bayesian inference with variational regularization is given by *Gibbs distributions*:

$$p_{pr}(u) \propto \exp(-\lambda \mathcal{J}(u)) \quad (7)$$

Here, $\mathcal{J}(u)$ is an energy functional penalizing unwanted features of u , and $\lambda > 0$ is a scaling parameter that is called the *regularization parameter*. Now, after suppressing terms not dependent on u , the MAP estimate is given by

$$\begin{aligned} \hat{u}_{\text{MAP}} &= \operatorname{argmax}_{u \in \mathbb{R}^n} \left\{ \exp \left(-\frac{1}{2\sigma^2} \|m - Au\|_2^2 - \lambda \mathcal{J}(u) \right) \right\} \\ &= \operatorname{argmin}_{u \in \mathbb{R}^n} \{ \|m - Au\|_2^2 + (2\sigma^2\lambda) \mathcal{J}(u) \} \end{aligned} \quad (8)$$

This is a *Tikhonov-type* regularization of equation (1) [17].

In this article, we only consider Gibbs priors with a L1 norm type energy functional:

$$\mathcal{J}(u) = |Du|, \quad (9)$$

where $D \in \mathbb{R}^{l \times n}$, and $|\cdot|$ denotes the L1 norm in \mathbb{R}^l . Although such priors may seem like a generic extension of the one dimensional *Laplace distribution* to a multivariate setting, we note here that multivariate generalizations of Laplace distributions are commonly defined in a different way (see, e.g. [16]). Concrete examples of L1-type priors will be given in Section 3. For the methods presented here, we require $l \leq n$, $\operatorname{rank}(D) = l$ and $\ker(D) \cap \ker(A) = 0$. In forthcoming work, we will extend the sampler proposed in Section 2.3 to more general settings.

Remark: For the sake of an intuitive presentation of the Bayesian formulation of inverse problems, we started from a deterministic setting. For a more detailed and rigorous description on how to derive a discrete, computational model of the continuous inverse problem starting in the Bayesian framework, we refer to [35].

2.2. Posterior Inference using MCMC Sampling

General Principles: In typical inverse problems scenarios, the dimension n of the unknowns u is very large (in our computational examples, we will study a scenario where the limit $n \rightarrow \infty$ is of central interest). Therefore, the integration to compute the CM estimate (6) is intractable by means of traditional quadratures. Interval, conditional covariance and histogram estimates and even more sophisticated topics in Bayesian inference like *marginalization*, *model selection* or *experiment design* [52] also rely on integration tasks and can, thus, not be computed by such an approach as well. Integration by *Monte Carlo* methods can avoid these difficulties. A sequence of points $u_i, i = 1, \dots, K$ is constructed, which is distributed like the posterior (the construction schemes are called *sampler* or *sampling methods*). If they were drawn independently, the *law of large numbers* would guarantee that

$$\frac{1}{K} \sum_{i=1}^K f(u_i) \xrightarrow{K \rightarrow \infty} \mathbb{E}[f(u)|m] = \int_{\mathbb{R}^n} f(u) p_{\text{post}}(u|m) du \quad (10)$$

for any measurable f almost surely and in L1 with rate $O(K^{-1/2})$. This means that the empirical mean of the sequence $f(u_i), i = 1, \dots, K$ converges to the expected value of $f(u)$ w.r.t the posterior [32]. A difficulty in our setting is that the posterior is not given in a form that allows for drawing independent samples. It is only known up to a normalizing constant (the model-evidence) and does not belong to a class of distributions for which independent sampling schemes are known. However, by the *strong ergodic theorem*, the above convergence (and its rate) still holds if the sequence is dependent, but originates from an *ergodic Markov chain* that has $p_{\text{post}}(u|m)$ as its *equilibrium distribution* [32]. Techniques to construct such chains are called *Markov chain Monte Carlo (MCMC)* methods. A huge number of different MCMC methods have been proposed. However, no method is known, which exhibits a good performance for all types of distributions. For a comprehensive overview, we refer to [37], for the application to inverse problems, see [31]. Most MCMC methods rely on one of two basic sampling schemes, which we will introduce in the next sections. Instead of comparing all possible and sophisticated variants of these schemes in our studies, we will use a small number of simple variants and focus on the differences between the two basic schemes for L1-type priors.

Metropolis-Hastings Sampling: For the ease of presentation, we denote the target probability density we want to sample by $p(x), x \in \mathbb{R}^n$. The *Metropolis-Hastings (MH)* algorithm [39, 25] is a very simple rule to generate a Markov chain:

Algorithm 1. (Metropolis-Hastings Sampling) Let $q(x, y) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a function satisfying $\int q(x, y) dy = 1$ for all $x \in \mathbb{R}^n$ (proposal distribution) and $x_0 \in \mathbb{R}^n$ an initial state. Define burn-in size K_0 and sample size K .

For $i = 1, \dots, K_0 + K$ do:

- 1 Draw y from the proposal distribution $q(x_{i-1}, y)$.

2 Compute the acceptance ratio

$$r(x_{i-1}, y) = \min \left(1, \frac{p(y) q(y, x_{i-1})}{p(x_{i-1}) q(x_{i-1}, y)} \right).$$

3 Draw $\theta \in [0, 1]$ from a uniform probability density.

4 If $r \geq \theta$, set $x_i = y$, else set $x_i = x_{i-1}$.

Return x_{K_0+1}, \dots, x_K .

Note that the restrictions on $p(x)$ for this scheme are minimal: We only have to know $p(x)$ up to a scaling factor, as only ratios of probabilities are used, and we only need to be able to evaluate $p(x)$ for any given x . Each sampling step requires one such evaluation (in inverse problems, the computational demanding part of this evaluation is usually applying the forward mapping A). The numerical implementation of the raw MH scheme is trivial. MH can, thus, be considered as a “black-box sampler”, which explains its success in many different application areas [37].

However, while the scheme works for all kinds of proposal distributions in theory, its application is only feasible if $q(x, y)$ leads to a chain that moves “fast” in the sampling space with respect to computational speed. This way, the important regions of the sampling space are explored reasonably fast and consecutive samples are as uncorrelated as possible (which improves the convergence in (10)). These requirements are hard to fulfill in practice. We refer to the discussions in [31, 37]. Usually, one ends up in a well-known dilemma of tuning different opposing parameters by manual inspection of different chain characteristics (see Section 2.4). Additionally, different applications usually require to develop and implement specific proposal distributions. As a consequence, a huge number of different MH-based schemes exist [37]. Especially for inverse problems, sophisticated algorithms that include automatic tuning procedures for the sampling parameters have been developed [22, 23, 9]. However, as mentioned earlier, a detailed comparison of their performance in sparse Bayesian inversion is not the topic of this publication. We rather want to compare the basic variants of MH algorithms and their performance to basic variants of Gibbs sampling algorithms. Therefore, we will use three proposal distributions that are commonly applied in practice because of their simplicity. They all belong to the class of *symmetric random-walk* Metropolis schemes [37]:

$$y = x + \vartheta, \quad \mathbb{E}(\vartheta) = 0, \quad p_\vartheta(\omega) \propto g(\|\omega\|_2) \forall \omega \in \mathbb{R}^n, \quad (11)$$

for a suitable, non-negative function g . This means that a new proposal y is generated by perturbing the current state x in a random, unbiased, symmetric way. Thus, $q(x, y) \propto g(\|x - y\|)$, and q vanishes from the acceptance ratio (1). The three choices of ϑ we will use are:

- 1) **MH-Iso:** All components of x are updated: $\vartheta_i \sim \mathcal{N}(0, \kappa^2)$, $\forall i$.
- 2) **MH-Ncom:** $1 < n^* < n$ components i_1, \dots, i_{n^*} of x are randomly chosen and are updated while the other components remain unchanged: $\vartheta_i \sim \mathcal{N}(0, \kappa^2)$. if $i \in \{i_1, \dots, i_{n^*}\}$, else $\vartheta_i = 0$.
- 3) **MH-Si:** One component i_* of x is randomly chosen and updated while all other components remain unchanged: $\vartheta_{i_*} \sim \mathcal{N}(0, \kappa^2)$, $\vartheta_{[-i_*]} = 0$.

Here, $\vartheta_{[-i]}$ denotes all components of ϑ except the i^{th} one. The concrete choice of n^* and κ will be explained in Section 3.

Gibbs Sampling: In certain scenarios, direct sampling of a n -dim multivariate distribution is not possible or computationally too expensive, but direct sampling from conditioned (thus, lower dimensional) versions of that distribution is feasible. In such a situation, *Gibbs sampling* can be applied. By successive sampling from the lower dimensional conditional distributions while changing the coordinates, which are fixed in each step, a Markov chain is generated [20, 18]. The most basic scheme is given by:

Algorithm 2. (Single Component Gibbs Sampling) Let $x_0 \in \mathbb{R}^n$ an initial state. Define burn-in size K_0 and sample size K
For $i = 1, \dots, K_0 + K$ do:

- Set $x_i := x_{i-1}$.

For $j = 1, \dots, n$ do:

- 1 Set $s = j$ (systematic scan) or draw s randomly from $\{1, \dots, n\}$ (random scan).
- 2 Draw $(x_i)_s$ from the conditional, 1-dim density $p(\cdot | (x_i)_{[-s]})$.

Return x_{K_0+1}, \dots, x_K .

We will abbreviate the systematic scan version of the above sampler as **SysGibbs** and the random version (which requires the extra computational effort of picking a random coordinate) as **RnGibbs**.

The basic Gibbs sampling scheme can be very slow if the correlations between the single components x_i are strong. This occurs naturally in typical under-determined inverse problems. In this case, the conditional distributions differ considerably from the corresponding marginal ones. As a consequence, the chain moves very randomly, exploring the search space very slowly. To address this problem, *overrelaxed* variants of Gibbs sampling have been proposed. The idea behind them are similar to those used in overrelaxation techniques for the iterative solution of systems of linear equations [48]. The specific form of overrelaxation that we will apply and examine was proposed in [43], and relies on *order statistics*. Step 2 in Algorithm 2 is replaced by:

Algorithm 3. (Ordered Overrelaxation)

- 2.1 Draw N_O random values from the conditional, 1-dim density $p(\cdot | (x_i)_{[-s]})$, where $N_O \in \mathbb{N}$ is odd.
- 2.2 Arrange these N_O values plus the old value $(x_i)_s$ in non-decreasing order, labeling them as follows:

$$(x_i)_s^{(0)} \leq (x_i)_s^{(1)} \leq \dots \leq (x_i)_s^{(t)} = (x_i)_s \leq \dots \leq (x_i)_s^{(N_O)} \quad (12)$$

- 2.3 Replace $(x_i)_s$ by $(x_i)_s^{(N_O-t)}$.

The value of N_O functions like an overrelaxation parameter. The larger the value of N_O , the larger the effect of overrelaxation and more randomness of the sampling process is suppressed. For symmetric densities, the current value of the component is mirrored at the mean and the whole chain moves on an iso-probability level of the density in the limit of $N_O \rightarrow \infty$. We will discuss more details of ordered overrelaxation in Section 3.1.5. We will denote the overrelaxed versions of SysGibbs and RnGibbs by appending " ON_O ", e.g., "SysGibbsO7" denotes the systematic scan Gibbs Sampler with ordered overrelaxation using $N_O = 7$. While the basic scheme for ordered overrelaxation requires N_O times more computation time compared to Algorithm 3, an efficient implementation is given in [43] that renders the computation time nearly independent of N_O . We will present this form after the next paragraph.

2.3. Implementation of Gibbs Sampling for L1-type Priors

In this section, the main contributions of this article are presented. The general Gibbs sampling schemes (Algorithms 2 and 3) need to be implemented in an efficient way. For this, we will first derive a way to compute a simple representation of the conditional single component density and then explain how to implement an exact, explicit and numerically robust sampler for it.

Conditional Densities for L1-type Priors: We will now derive the single component conditional densities required by Algorithms 2 and 3 for our setting (cf. Section 2.1). Because $\text{rank}(D) = l$ (cf. Section 2.1), we can find $v_1, \dots, v_l \in \mathbb{R}^n$ such that $D v_i = e_i$, (where e_i denotes the i^{th} unit vector in \mathbb{R}^l) and $v_{l+1}, \dots, v_n \in \mathbb{R}^n$ such that v_1, \dots, v_n form a basis of \mathbb{R}^n . Then, we have

$$u = \sum_{i=1}^n \xi_i v_i, \quad D u = (\xi_1, \dots, \xi_l)^t \quad \text{and} \quad |D u| = \sum_{i=1}^l |\xi_i|.$$

With $V := [v_1, \dots, v_n]$, we can transform the posterior to:

$$\begin{aligned} \exp \left(-\frac{1}{2\sigma^2} \|m - A u\|_2^2 - \lambda |D u| \right) &= \exp \left(-\frac{1}{2\sigma^2} \|m - A V \xi\|_2^2 - \lambda \sum_{i=1}^l |\xi_i| \right) \\ &:= \exp \left(-\|\bar{m} - \Psi \xi\|_2^2 - \lambda \sum_{i=1}^l |\xi_i| \right), \end{aligned} \quad (13)$$

where $\bar{m} := m/(\sqrt{2}\sigma)$ and $\Psi := (A V)/(\sqrt{2}\sigma)$. Because the transformations are linear, no specific attention to the correct transformation of probability densities has to be paid. Now let ψ_i be the i^{th} column of Ψ , $\Psi_{[-i]}$ be Ψ without the i^{th} column and $\xi_{[-i]}$ be ξ without the i^{th} entry. Then

$$\bar{m} - \Psi \xi = (\bar{m} - \Psi_{[-i]} \xi_{[-i]}) - \psi_i \xi_i := \varphi_{[-i]} - \psi_i \xi_i. \quad (14)$$

Consider the conditional posterior of ξ_i given \bar{m} and $\xi_{[-i]}$:

$$\begin{aligned} p(\xi_i | \bar{m}, \xi_{[-i]}) &\propto \exp \left(-\|\varphi_{[-i]} - \psi_i \xi_i\|_2^2 - \lambda |\xi_i| \cdot \mathbf{1}_{\{i \leq l\}} \right) \\ &= \exp \left(-\langle \varphi_{[-i]} - \psi_i \xi_i, \varphi_{[-i]} - \psi_i \xi_i \rangle - \lambda |\xi_i| \cdot \mathbf{1}_{\{i \leq l\}} \right) \\ &\propto \exp \left(-\|\psi_i\|_2^2 \xi_i^2 + 2\psi_i^t \varphi_{[-i]} \xi_i - \lambda |\xi_i| \cdot \mathbf{1}_{\{i \leq l\}} \right) \end{aligned} \quad (15)$$

To ease the following presentation, we define:

$$x := \xi_i; \quad a := \|\psi_i\|_2^2; \quad b := 2\psi_i^t \varphi_{[-i]} = 2 \left[\psi_i^t \bar{m} - (\psi_i^t \Psi_{[-i]}) \xi_{[-i]} \right]; \quad c := \lambda \cdot \mathbf{1}_{\{i \leq l\}} \quad (16)$$

Thus, the problem of sampling from the single component conditional densities is reduced to sampling from the 1-dim density

$$p(x) \propto \exp(-a x^2 + b x - c |x|), \quad (17)$$

once a , b and c have been computed by (16). In the next paragraph, we will describe how to use the *inverse cumulative distribution method* [31] to sample from (17). Concerning the practical implementation of computing a , b and c in a fast way, note that a and c can be precomputed and only b depends on the current state of the chain ξ through the term $(\psi_i^t \Psi_{[-i]}) \xi_{[-i]}$. The most efficient way to compute this term strongly depends on the form of A and V , on the problem size n and on the hardware

available. If enough working memory is available to store the $n \times n$ matrix $\Phi := \Psi^t \Psi$, the most efficient way is to compute

$$(\psi_i^t \Psi_{[-i]}) \xi_{[-i]} = \xi^t \Phi_{(\cdot, i)} - \xi_i \|\psi_i\|_2^2, \quad (18)$$

because the most extensive operation is a scalar product of dimension n . In the scenario examined in [35], A is a symmetric convolution operator and V an inverse wavelet transform, i.e., v_j are the wavelets. For large k and n (as encountered, e.g., in 2D or 3D imaging applications), it is infeasible to compute and store the matrix form of A , V or Ψ . Then, it is advantageous to use

$$\begin{aligned} (\psi_i^t \Psi_{[-i]}) \xi_{[-i]} &= \psi_i^t (\Psi \xi) - \xi_i \|\psi_i\|_2^2 = \frac{1}{2\sigma^2} (A \cdot V e_i)^t (A V \xi) - \xi_i \|\psi_i\|_2^2 \\ &= \frac{1}{2\sigma^2} [A \cdot (V e_i)]^t [A (V \xi)] - \xi_i \|\psi_i\|_2^2 = \frac{1}{2\sigma^2} v_i^t [(A^t A) \cdot (V \xi)] - \xi_i \|\psi_i\|_2^2. \end{aligned} \quad (19)$$

Here, $(V \xi)$ can be realized using the *fast wavelet transform*, while the double convolution by $(A^t A)$ can be substituted by a single convolution with a different kernel and realized by the *fast Fourier transform*.

Explicit 1D Sampling: Sampling from continuous 1-dim distributions by the inverse cumulative distribution method follows a simple rule: Let $F(y) := \int_{-\infty}^y p(x) dx$ be the *cumulative distribution function (cdf)* and r be a random number uniformly drawn from $[0, 1]$. Then, $y = F^{-1}(r)$ is distributed like $p(x)$ (see [37]). We can also use this concept to provide an equivalent implementation of Algorithm 3. For for a given N_O :

Algorithm 4. (CDF Implementation of Ordered Overrelaxation)

- 2.1 Compute $r = F[(x_i)_s]$, which lies in $[0, 1]$.
- 2.2 Let r' be the random ordered overrelaxation of r w.r.t to the uniform distribution on $[0, 1]$ and N_O (computed with Algorithm 3).
- 2.3 Replace $(x_i)_s$ by $F^{-1}(r')$.

For more details, we refer to [43]. Turning these rules into efficient sampling schemes requires a fast and stable way to invert F , which is defined by an integral. Using numerical integration for this purpose often fails to render fast and robust sampling algorithms. In this paragraph, we will present a scheme that relies on the *inverse complementary error function* (erfcinv) for which efficient and stable implementations are known.

First, we compute the normalization factor for $p(x) \propto \exp(-a x^2 + b x - c |x|)$. Splitting the integral from $-\infty$ to ∞ into two parts (from $-\infty$ to 0 and the rest) yields subproblems that can be treated like the normalization of the normal distribution (completing the square and a linear integral transformation). This leads to:

$$\begin{aligned} \mathcal{N} &:= \int_{-\infty}^{\infty} \exp(-a x^2 + b x - c |x|) dx \\ &= \frac{1}{2} \sqrt{\frac{\pi}{a}} \left[e^{\frac{(b+c)^2}{4a}} \text{erfc}\left(\frac{b+c}{2\sqrt{a}}\right) + e^{\frac{(c-b)^2}{4a}} \text{erfc}\left(\frac{c-b}{2\sqrt{a}}\right) \right] \\ &:= \chi [\tilde{e}_+ \text{erfc}(\alpha_+) + \tilde{e}_- \text{erfc}(\alpha_-)], \end{aligned} \quad (20)$$

where $\text{erfc}(y) := \frac{2}{\sqrt{\pi}} \int_y^{\infty} e^{-t^2} dt$ denotes the complementary error function. The cdf is given by:

$$\text{cdf}(y) := \frac{1}{\mathcal{N}} \int_{-\infty}^y \exp(-a x^2 + b x - c |x|) dx$$

$$= \frac{\chi}{\mathcal{N}} \cdot \begin{cases} \tilde{e}_+ \operatorname{erfc}(-\sqrt{a}y + \alpha_+) & , \text{ if } y < 0, \\ \tilde{e}_+ \operatorname{erfc}(\alpha_+) + \tilde{e}_- [\operatorname{erfc}(\alpha_-) - \operatorname{erfc}(\sqrt{a}y + \alpha_-)] & , \text{ if } y > 0. \end{cases} \quad (21)$$

Inverting this cdf for a given $r \in [0, 1]$ is simple. To find $y = \operatorname{cdfinv}(r)$ we first check if $y < 0$ by using the cdf for this domain. Let

$$\begin{aligned} z := \operatorname{erfcinv}\left(\frac{r\mathcal{N}}{\chi\tilde{e}_+}\right) &= \operatorname{erfcinv}\left\{\frac{r\chi[\tilde{e}_+ \operatorname{erfc}(\alpha_+) + \tilde{e}_- \operatorname{erfc}(\alpha_-)]}{\chi\tilde{e}_+}\right\} \\ &= \operatorname{erfcinv}\left\{r\left[\operatorname{erfc}(\alpha_+) + \frac{\tilde{e}_-}{\tilde{e}_+} \operatorname{erfc}(\alpha_-)\right]\right\} \\ &= \operatorname{erfcinv}\left\{r\left[\operatorname{erfc}(\alpha_+) + \exp\left(-\frac{bc}{a}\right) \operatorname{erfc}(\alpha_-)\right]\right\}, \end{aligned} \quad (22)$$

then, y is given by $y = -(z - \alpha_+)/\sqrt{a}$. If it turns out that this y fulfills $y > 0$, the other half of the cdf has to be inverted. Let

$$\begin{aligned} z &:= \operatorname{erfcinv}\left\{\left[-\frac{r\mathcal{N}}{\chi} + \tilde{e}_+ \operatorname{erfc}(\alpha_+) + \tilde{e}_- \operatorname{erfc}(\alpha_-)\right] \tilde{e}_-^{-1}\right\} \\ &= \operatorname{erfcinv}\left\{(1-r)\left[\exp\left(\frac{bc}{a}\right) \operatorname{erfc}(\alpha_+) + \operatorname{erfc}(\alpha_-)\right]\right\}. \end{aligned} \quad (23)$$

Then, y is given by $y = (z - \alpha_-)/\sqrt{a}$.

The complementary error function and its inverse are difficult to handle numerically, because there are no identities that allow to rescale or shift their evaluation to other intervals. Therefore, a robust numerical implementation of formulas (21), (22) and (23) is rather involved. For the sake of a concise presentation, we present all details in Appendix B.

2.4. MCMC Convergence Diagnostics

Assessing the efficiency of a sampling algorithm for a general purpose rather than a specific aim is a difficult task [37]. Two types of *convergence diagnostics* are usually applied: *Qualitative diagnostics* rely on the visual inspection of some property of the chain $u_i, i = 1, \dots, K$. In contrast, *quantitative diagnostics* try to compute characteristics that can be used to guide the sampling algorithm in an automated fashion. This should allow unexperienced users to perform “black box” Bayesian inference. Despite a lot of research on these topics [12, 4, 46, 51], no universal method is known. For our purpose, a qualitative autocorrelation analysis is appropriate. For a test function $g : \mathbb{R}^n \rightarrow \mathbb{R}^1$, the *autocorrelation function (acf)* $R : \{0, \dots, K-1\} \rightarrow [-1, 1]$ of the series $g_i := g(u_i), i = 1, \dots, K$ is given by:

$$R(\tau) := \frac{1}{(K-\tau)\hat{\varrho}} \sum_{i=1}^{K-\tau} (g_i - \hat{\mu})(g_{i+\tau} - \hat{\mu}) \quad (24)$$

$$\hat{\varrho} := \frac{1}{K} \sum_{i=1}^K (g_i - \hat{\mu})^2, \quad \hat{\mu} := \frac{1}{K} \sum_{i=1}^K g_i \quad (25)$$

(Note that there are other possibilities to define R , but we need $R(0) = 1$). The value of $R(\tau)$ is referred to as the *lag- τ autocorrelation* w.r.t. g . A fast decrease of the acf indicates that consecutive samples get mutually independent quite soon (if the u_i would be independent, then, $R(\tau) = \delta_{(\tau,0)}$). For practical considerations, the decrease of autocorrelation w.r.t. to the raw number of samples drawn is not decisive if different

samplers are compared. A method that has a slower decrease than others might still outperform them if it produces new samples considerably faster. In such situations, one would subsample the chain to get rid of highly correlated samples and to save memory. Note that the notion of "one" sample is quite arbitrary anyway. In the SysGibbs sampler, one speaks of a "new" sample, if all components of u are updated, in the MH-Si sampler one speaks of a "new" sample, if one component is updated. To address this, we will normally scale the acf by the computation time per sample t_s : $R^*(t) := R(t/t_s)$ for all $t = i \cdot t_s, i \in \{0, \dots, K-1\}$, if we compare conceptually different sampling methods. $R^*(t)$ measures how fast a sampler can produce a certain loss in autocorrelation, which is of main interest for practical applications. However, while $R^*(t)$ is more decisive to compare different samplers, it relies on their concrete implementation[‡].

Normally, the test function g is chosen with respect to the specific aim of inference. For instance, one could use the distance to the empirical mean of the whole chain if CM estimation is performed, or the projection onto a specific coordinate if that coordinate should be marginalized. Then, the rate of autocorrelation decrease is a measure of the efficiency of the chain for the specific inference aim. For our general purpose, we will test the "worst case". We project onto the direction of the largest variance, i.e., the first eigenvector ν_1 of the covariance matrix C of the posterior:

$$g(u_i) := \langle \nu_1, u_i \rangle \quad (26)$$

In general, the chain should have most problems to reduce the correlation of subsequent samples in this direction ν_1 . For each scenario we examine, the covariance matrix C of the posterior is estimated from a long (sub-sampled) chain of the RnGibbs sampler, as this sampler will turn out to be the most reliable at a high performance. Note that this choice does not give an advantage to the RnGibbs sampler in the autocorrelation analysis but rather a disadvantage if the other samplers would have other directions of highest variance. We checked that this is not the case in a test scenario we examined in preliminary studies.

Other possible MCMC convergence diagnostic plots that are commonly used are plots of $\log[p(u_i|m)]$ or of single components $(u_i)_j$. Such plots are good to detect possible *multimodality* of the posterior and to determine a sufficient number of burn in steps K_0 . Multimodality is not an issue in our case, as the posterior is log-concave (the energy $-\log[p(u|m)]$ is convex). The burn-in length is an important factor for the practicability of the algorithms (and we will address this issue in our studies) but it is a difficult measure for a fair and definite comparison of the sampling methods. First, it crucially relies on the initialization of the chain, so one would have to compare all methods for various common initialization strategies, which is not really feasible and too application specific. Second, for the Metropolis-Hastings schemes, an adaptation of the sampling parameters to is usually carried out in the burn-in phase with the aim to optimize the performance of the chain in the real run. We will introduce this topic in Section 3.1.2. The consequence is that K_0 also depends on the adaptation scheme, which renders the problem of a meaningful comparison even worse.

[‡] We implemented all samplers in Matlab and optimized them to yield the best possible performance. As mentioned in Section 1.2, we originally intended to use the MH samplers in the scenario examined in Section 3.1. As their results were unsatisfactory even after a careful optimization of their implementation, we decided to develop the Gibbs samplers presented in this paper.

3. Results

In this section, we compare the sampling algorithms for two scenarios: Edge-preserving, TV-based image deblurring in 1D and impulse prior based image deblurring in 2D. All algorithms have been implemented in Matlab and have been optimized to the best possible performance. All results have been computed on the same CPU architecture limiting Matlab to a single computational thread, i.e., to use a single CPU core with 2.80GHz (parallelization is discussed in Section 4). We paid special attention that the computation times are as comparable as possible.

3.1. Edge-Preserving Bayesian Inversion in 1D

A popular case of L1-type priors arises from *edge-preserving* image reconstruction. The task is to reconstruct a spatially distributed intensity image that is known to consist of piecewise homogeneous parts with sharp edges from indirect, noisy measurements (e.g., the recovery of the body's organs and their boundaries from X-ray computed tomography data [34, 49, 31]). Using Gaussian, i.e., L2 -type priors smooths the image edges in such situations. In contrast, *total variation* (*TV*) priors, which rely on the L1 norm of the first spatial derivatives, are able to retain them [47, 31, 38, 6]. The use of TV priors in Bayesian inference has led to interesting theoretical questions. It was discovered that it is not possible to formulate the conventional TV prior in a *discretization invariant* way [36, 35], i.e., that the posterior converges to a well defined limit probability density when the level of discretization is increased while reflecting the a priori information of edge-preservation at all levels of discretization. If the TV prior is formulated such that it converges, it converges to a Gaussian smoothness prior, and, thus, the edge-preservation property is lost. This motivated research on whether and how it is possible to formulate edge-preservation as a priori information in a consistent, discretization invariant way in the Bayesian framework. Recently, *Besov space priors* have been proposed, which rely on a weighted L1 norm of wavelet basis coefficients [35, 33]. Such priors are L1-type priors with an invertible D , thus, posterior sampling by means of our Gibbs sampling algorithms can be performed with ease. In addition, modifications of the standard TV prior [11] and *hierarchical Bayesian models* [7, 8, 27, 26, 2] have been proposed for discretization invariant edge-preserving image reconstruction as well.

To address the problem of discretization invariance of a prior, one can, e.g., study the convergence of the corresponding CM estimate for $n \rightarrow \infty$. The problems of using MH-based samplers for CM estimation in high dimensions have already been noticed in [36, 33] and our research on alternative samplers has been motivated by these problems as well.

3.1.1. Setting We rely on the setting used in [36]. The motivation is to mimic a measurement made by a *charge coupled device* (*CCD*) used in digital cameras or medical imaging devices. These devices integrate the amount of light illuminating a certain pixel over a certain period of time. In the continuous model setting, we represent the unknown light intensity by a positive function $\tilde{u} : [0, 1] \rightarrow \mathbb{R}$, and the k pixels of the CCD device as a equidistant division of the subinterval $[\frac{1}{k+2}, \frac{k+1}{k+2}] \subset [0, 1]$, i.e., the j -th pixel is represented by the interval $[\frac{j}{k+2}, \frac{j+1}{k+2}]$. The measurement at the

j^{th} pixel is then given by:

$$m_j = \int_{\frac{j}{k+2}}^{\frac{j+1}{k+2}} \tilde{u}(t) dt + \varepsilon_j \quad (27)$$

For convenience, we will choose $k = 2^{L_m} - 2$ and $L_m = 5$. For discretizing \tilde{u} , we choose the grid $\{\frac{1}{n+1}, \dots, \frac{n}{n+1}\} \subset [0, 1]$, and let $n = 2^{L_u} - 1$ with $L_u > L_m$. The discretization of the forward mapping implied by (27) in terms of the $k \times n$ matrix A can then be implemented by the trapezoidal quadrature rule. The j^{th} row of A is given by

$$A_{(j,\cdot)} := [\underbrace{0, 0, \dots, 0}_{j2^{(L_u-L_m)}-1}, \underbrace{\frac{1}{2}h, h, h, \dots, h}_{2^{(L_u-L_m)}-1}, \frac{1}{2}h, 0, 0, \dots, 0], \quad (28)$$

where $h := \frac{1}{n+1}$ defines the grid size. The discrete TV prior with Neumann boundary conditions in our situation is given by:

$$p(u) \propto \exp \left(-\lambda_n \sum_{i=1}^{n-1} |u_{i+1} - u_i| \right) = \exp(-\lambda_n |Du|) \quad (29)$$

$$(30)$$

where $D \in \mathbb{R}^{(n-1) \times n}$ is given by $(Du)_i := u_{i+1} - u_i$, $i = 1, \dots, n-1$. We indexed λ_n by n to stress that we may choose it depending on the discretization level.

For the Gibbs sampler (cf. Sections 2.1 and 2.3), we note that $l = \text{rank}(D) = n-1$, and $v_1, \dots, v_{n-1} \in \mathbb{R}^n$ are given by step functions: $(v_i)_j = \mathbf{1}_{\{j > i\}}$. These are completed to a basis of \mathbb{R}^n by $(v_n)_j = 1 \forall j$. If we reorder them and define $V := [v_n, v_1, \dots, v_{n-1}]$, we can write V as

$$V_{(i,j)} = \begin{cases} 1 & \text{if } i \geq j \\ 0 & \text{else} \end{cases} \quad (31)$$

The unknown function \tilde{u} we actually use is the indicator function on $[\frac{1}{3}, \frac{2}{3}]$, see Figure 1(a). Measurement data m is generated using formula (27), see Figure 1(b). The standard deviation of the measurement noise σ is 0.001. We will examine different




Figure 1. Left: The unknown function $\tilde{u}(t)$. Right: The measurement data m .

combinations of n and λ_n :

- A: $n = 63$ in combination with $\lambda_n = 100, 200$ and 400 , respectively. Here, we focus on increasing the impact of the prior. The posterior will become less Gaussian because the weight of the L1-type TV prior is increased.

B: $n = 2^{L_u} - 1$ for $L_u = 7, 8, \dots$ with $\lambda_n = 25 \cdot \sqrt{n+1}$. With this scaling of λ_n , the posterior $p(u|m)$ converges for $n \rightarrow \infty$, but the edge-preserving property of the TV prior is lost, see [36] for details. The CM estimate u_{CM} converges to a smooth limit function, which will facilitate the visual validation of the results of the different MCMC methods.

3.1.2. Preliminaries

Choice of Parameters: For the MH schemes, the proper tuning of κ is essential. If it is very small, the proposals will always be accepted since the distribution to sample from is continuous. However, in return the exploration of the sampling space is slow. On the contrary, if it is too large, the differences in probability will be huge because the distribution is log-concave and new proposals will hardly be accepted. A good overview on this topic is given in [45, 42]. The remarkable result is that in high dimensions, having a total acceptance rate of new proposals of about 0.234 leads to an optimal efficiency independent of the distribution to sample from. Furthermore, this optimal efficiency hardly drops in the range between 0.1 and 0.4 of acceptance rate. This yields an easy to implement rule to tune κ : One could find the optimal κ in a preliminary MH-MCMC run and initialize the real MH-MCMC run with it. However, it turns out that this κ is only optimal once the chain has reached the main support of the distribution while it can hinder the chain from ever getting there (the burn-in length increases dramatically, see Section 2.4). For these reasons, on-line adaptation of κ is usually used. The empirical acceptance rate is monitored, and κ is increased if it is too high while κ is decreased if it is too low. The scheme we use is that every 10 000 samples, the empirical acceptance rate is computed and if it is above 0.35, κ is multiplied by 1.2 while it is multiplied by 0.8 if it is below 0.15. In theory, the resulting chain will then not be a Markov chain anymore (but it is still ergodic). However, in practice, using this scheme, κ hardly ever changes once the burn-in time is over and so the real chain is not affected.

For MH-Ncom, we have to choose n^* , i.e., the number of components that are updated in one step. We choose $n^* = \lfloor n^{7/12} \rfloor$, which roughly corresponds to the values used in [36].

Burn-in Times: As noted in Section 2.4, the sufficient amount of burn-in steps that have to be drawn can be deduced from observing $\log[p(u_i|m)]$. Once it oscillates around a constant value, the stationary part of the distribution is reached. Averaging $\log[p(u_i|m)]$ over a large number of independent chains that all started at the same initialization ($u \equiv 0$ in our case) removes the oscillations and allows to determine K_0 in an easy fashion. See Figure 2 for an example of such a plot. In Table 1, the burn-in steps K_0 and the corresponding computation times t_0 are listed for the combinations of n and λ that are examined in detail. It gives a first impression of how the methods scale with n and λ_n , but as noted in Section 2.4 it does not allow for a fair and detailed comparison.

3.1.3. General Autocorrelation Analysis As explained in Section 2.4, we will rely on autocorrelation plots for the projection of the samples onto the direction of maximal covariance as qualitative measures of the efficiency of the sampling algorithms. Figure 3 shows the autocorrelation plots $R(\tau)$ for $n = 63$ and varying λ . In Figure 4, the




Figure 2. Plots of $\log[p(u_i|m)]$ for $n = 1023$, $\lambda = 800$, using the RnGibbs sampler. Red \bullet , $*$, $+$: three independent realizations. Blue \times : the average of 5000 independent realizations.

Table 1. Necessary burn-in steps K_0 and computation time t_0 in seconds for each method and combination of n and λ when starting from $u \equiv 0$. The values were found by observing the log posterior probability averaged over a large number of independent chains. Listed as (K_0, t_0)

Method	Model parameters (n, λ)		
	(63,100)	(63,200)	(63,400)
MH-Iso	(4e5,1.8e1)	(4e5,1.9e1)	(5e5,2.3e1)
MH-Ncom	(4e5,2.3e1)	(4e5,2.5e1)	(5e5,2.9e1)
MH-Si	(5e5,2.8e1)	(5e5,3.0e1)	(6e5,3.4e1)
RnGibbs	(200,0.5e0)	(200,0.5e0)	(200,0.4e0)
RnGibbsO3	(200,0.9e0)	(200,1.0e0)	(200,0.9e0)
RnGibbsO7	(200,1.0e0)	(200,1.0e0)	(200,0.9e0)
SysGibbs	(400,1.0e0)	(500,1.3e0)	(500,1.3e0)
SysGibbsO3	(200,0.9e0)	(500,2.2e0)	(500,2.2e0)
SysGibbsO7	(200,0.9e0)	(500,2.3e0)	(500,2.2e0)

Method	Model parameters (n, λ)			
	(127,280)	(255,400)	(511,560)	(1023,800)
MH-Iso	(7e5,3.4e1)	(4e6,2.1e2)	(3e7,1.9e3)	(2e8,1.7e4)
MH-Ncom	(7e5,4.2e1)	(4e6,2.6e2)	(3e7,2.5e3)	(2e8,2.3e4)
MH-Si	(8e5,4.4e1)	(4e6,2.3e2)	(3e7,1.9e3)	(2e8,1.5e4)
RnGibbs	(80,0.4e0)	(50,0.4e0)	(30,0.5e0)	(20,0.6e0)
RnGibbsO3	(80,0.7e0)	(50,0.9e0)	(30,1.1e0)	(20,1.4e0)
RnGibbsO7	(80,0.7e0)	(50,0.9e0)	(30,1.1e0)	(20,1.5e0)
SysGibbs	(150,0.7e0)	(100,0.9e0)	(150,2.9e0)	(150,5.3e0)
SysGibbsO3	(150,1.4e0)	(150,2.6e0)	(150,5.2e0)	(150,1.0e1)
SysGibbsO7	(150,1.4e0)	(150,2.7e0)	(200,6.8e0)	(200,1.4e1)

corresponding temporal autocorrelation plots $R^*(t)$ are shown. The comparison is split up into MH-based vs. normal Gibbs samplers and normal vs. overrelaxed Gibbs samplers to reduce the number of plots shown in one figure. The plots for RnGibbsO3, SysGibbsO3 and MH-Ncom were omitted for the same reason: The plots

for RnGibbsO3 and SysGibbsO3 lie between the plots of RnGibbs and RnGibbsO7 and SysGibbs and SysGibbsO7, respectively. The plots of MH-Ncomp look similar to the ones of MH-Iso and lie between MH-Iso and MH-Si. Figures 5 and 6 show the autocorrelation plots $R(\tau)$ for varying n and $\lambda_n = 25 \cdot \sqrt{n+1}$. In Figure 7, the temporal autocorrelation plots $R^*(t)$ corresponding to Figure 5 are shown. The plots for RnGibbsO3, SysGibbsO3 and MH-Ncom are, again, omitted. In addition, the plots for $l = 10$ are not shown as the trends of the autocorrelation functions for growing n are already clearly visible.

Table 2 lists the lag $\tau_{0.01}$ for which the autocorrelation $R(\tau)$ drops below 1 % for the first time and the corresponding computation time $t_{0.01} = \tau_{0.01} \cdot t_s$.

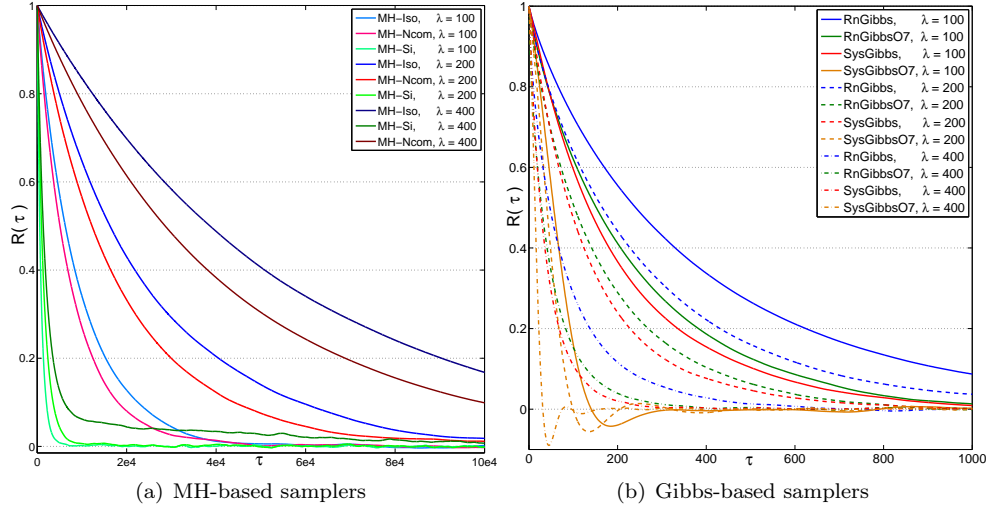


Figure 3. Autocorrelation plots $R(\tau)$ for $n = 63$, $\lambda_n = 100, 200$ and 400 .

3.1.4. Visual Results To get a visual impression of the sampling results, CM estimates are computed using the different samplers at different computation times for $n = 1023$, $\lambda_n = 25 \cdot \sqrt{n+1}$. The computation times examined are $t^* = 1$ s, 10 s, 1 minute, 1 hour and 1 day, respectively. Practically, a long chain with $K_0 = 0$ was generated and sub-chains corresponding to all the samples drawn before t^* were extracted. Then, CM estimates were computed from the sub-chains by discarding $K_0^* = \min(K_0, 0.5 \cdot K^*)$ burn-in samples, where K_0 are the burn-in steps listed in Table 1, and K^* denotes the number of samples in the subchain. The results are shown in Figure 8. We have to emphasize that we did not chose to show the CM estimate for the TV prior because the reconstruction is convincing. In fact, as explained in Section 3.1 they are extremely smooth compared to the corresponding MAP estimates and, thus, bad reconstructions of the discontinuous \tilde{u} . However, this smoothness is very useful for gaining a visual impression of the convergence and the properties of the different sampling schemes: The CM estimate computed from the chain converged once it is smooth. To demonstrate the capabilities of the newly developed Gibbs samplers for the practical use, we also examine the theoretical questions addressed in [36]. For the choice of $\lambda_n \propto \sqrt{n+1}$, the TV prior converges to a smoothness prior. To support this finding with numerical simulations the CM estimate was computed

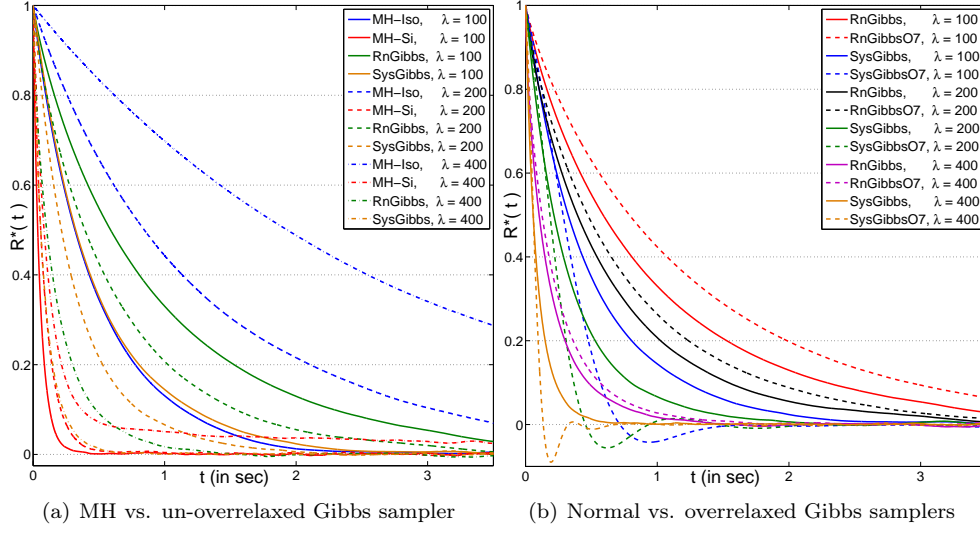


Figure 4. Temporal autocorrelation plots $R^*(t)$ for $n = 63$, $\lambda_n = 100, 200$ and 400 .




Figure 5. Autocorrelation plots $R(\tau)$ for varying n and λ for MH and un-overrelaxed Gibbs sampler. Note that the τ axis starts at $\tau = 1$ and is scaled logarithmically.

for $n = 63, 255, 1023, 4095$ in [36] using the MH-Ncom sampler. Although the whole computation took about a month of time on a desktop PC equipped with a 2.8 GHz single core CPU, the authors admitted that the results were only partly satisfying. In Figure 9, we show the CM estimate computed for $n = 63, 255, 1023, 4095, 16383, 65535$ using the RnGibbs sampler on a comparable CPU. Again, the CM estimates are only shown because the increasing smoothness of the CM estimates for growing n allows for the visual inspection of the chain convergence.




Figure 6. Autocorrelation plots $R(\tau)$ for varying n and λ for overrelaxed Gibbs sampler. Note that the τ axis starts at $\tau = 1$ and is scaled logarithmically.




Figure 7. Temporal autocorrelation plots $R^*(t)$ for varying n and λ for MH un-overrelaxed Gibbs sampler. Note that the t axis starts at the smallest t_s and is scaled logarithmically.

3.1.5. Normal vs. ORR Gibbs Sampling Using oriented overrelaxation removes a certain amount of randomness from the generated chains. This may lead to a faster exploration of the posterior distribution, but can also enhance non-ergodic tendencies of the sampling approaches. In Figure 10, more detailed autocorrelation plots for the 1D scenario using $n = 1023$, $\lambda = 800$ are shown. Whether oriented overrelaxation is practically advantageous relies on the computational cost of sampling from the single component density (17) costs compared to the other computation steps. Using oriented overrelaxation to sample from (17) takes roughly twice as much computation

Table 2. Lag of 1% auto correlation for each method and combination of n and λ in terms of (samples,computation time)

Method	Model parameters (n, λ)		
	(63,100)	(63,200)	(63,400)
MH-Iso	(4.1e4,2.1e0)	(1.2e5,6.2e0)	(2.1e5,1.1e1)
MH-Ncom	(4.2e4,2.6e0)	(1.0e5,6.4e0)	(2.1e5,1.0e1)
MH-Si	(4.7e3,0.3e0)	(7998,0.5e0)	(8.8e4,5.4e0)
RnGibbs	(1685,4.1e0)	(1402,3.3e0)	(561,1.2e0)
RnGibbsO3	(1239,5.8e0)	(983,4.6e0)	(395,1.8e0)
RnGibbsO7	(1056,5.5e0)	(811,3.7e0)	(318,1.4e0)
SysGibbs	(985,2.4e0)	(810,1.9e0)	(242,0.5e0)
SysGibbsO3	(412,1.9e0)	(242,1.1e0)	(76,0.3e0)
SysGibbsO7	(137,0.7e0)	(97,0.4e0)	(31,0.1e0)

Method	Model parameters (n, λ)			
	(127,280)	(255,400)	(511,560)	(1023,800)
MH-Iso	(1.1e6,5.0e1)	(4.6e6,2.5e2)	(2.7e7,1.9e3)	(1.3e8,1.3e4)
MH-Ncom	(9.4e5,5.4e1)	(3.2e6,2.2e2)	(2.1e7,1.8e3)	(1.8e8,2.3e4)
MH-Si	(4.8e4,3.1e0)	(2.1e6,1.2e2)	(3.1e7,2.1e3)	(2.9e8,2.5e4)
RnGibbs	(2017,9.2e0)	(1014,8.7e0)	(46,0.8e0)	(39,1.3e0)
RnGibbsO3	(1006,8.9e0)	(1052,2.0e1)	(31,1.1e0)	(29,2.1e0)
RnGibbsO7	(953,8.7e0)	(473,8.7e0)	(28,0.9e0)	(24,1.8e0)
SysGibbs	(770,3.4e0)	(270,2.3e0)	(9,0.1e0)	(12,0.4e0)
SysGibbsO3	(230,2.0e0)	(165,2.9e0)	(8,0.3e0)	(7,0.5e0)
SysGibbsO7	(126,1.2e0)	(153,2.8e0)	(7,0.2e0)	(6,0.4e0)

time as not using it, almost independent from N_O (when using Algorithm 4). If the other computation steps in the whole sampling scheme take way more time (i.e., the computation of b , see Section 2.3), this extra computational cost is negligible. In our scenario, the ratio between the total computation time per sample t_s for the RnGibbsO7 and the RnGibbs sampler varies considerably. It drops from 1.93 for the fast but memory consuming implementation (see Section 2.3) to compute b and $n = 63$ to 1.02 for the slower implementation to compute b and $n = 65\,535$.

3.2. Image Deblurring with Impulse Prior in 2D

3.2.1. Setting As a second example, we consider 2D image deblurring with a simple L1 prior, i.e., $D = I_n$ (also called *impulse prior*). The unknown intensity function $\tilde{u} : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ is shown in Figure 11(b). It consists of a couple of circular spots of constant intensity whose radii and intensities slightly vary between single spots. The forward mapping is given by a convolution with a Gaussian kernel with standard deviation of 0.015. Measurement data is generated by integrating the resulting convoluted image over 513×513 regular pixel and adding noise. The relative noise level is 0.1, i.e., the standard deviation σ of the measurement noise is 0.1 times the maximal intensity of the noiseless signal. The resulting measurement data is shown in Figure 11(c). The image will be reconstructed on the same pixel grid used for the measurement using Neumann boundary conditions, thus, the dimension of the unknowns n is $511^2 = 261\,121$. To avoid an inverse crime, the grid used for the generation of the measurement data was 4 times finer. For the MH samplers, the same on-line adaptation of κ was used as explained in Section 3.1.2. The update

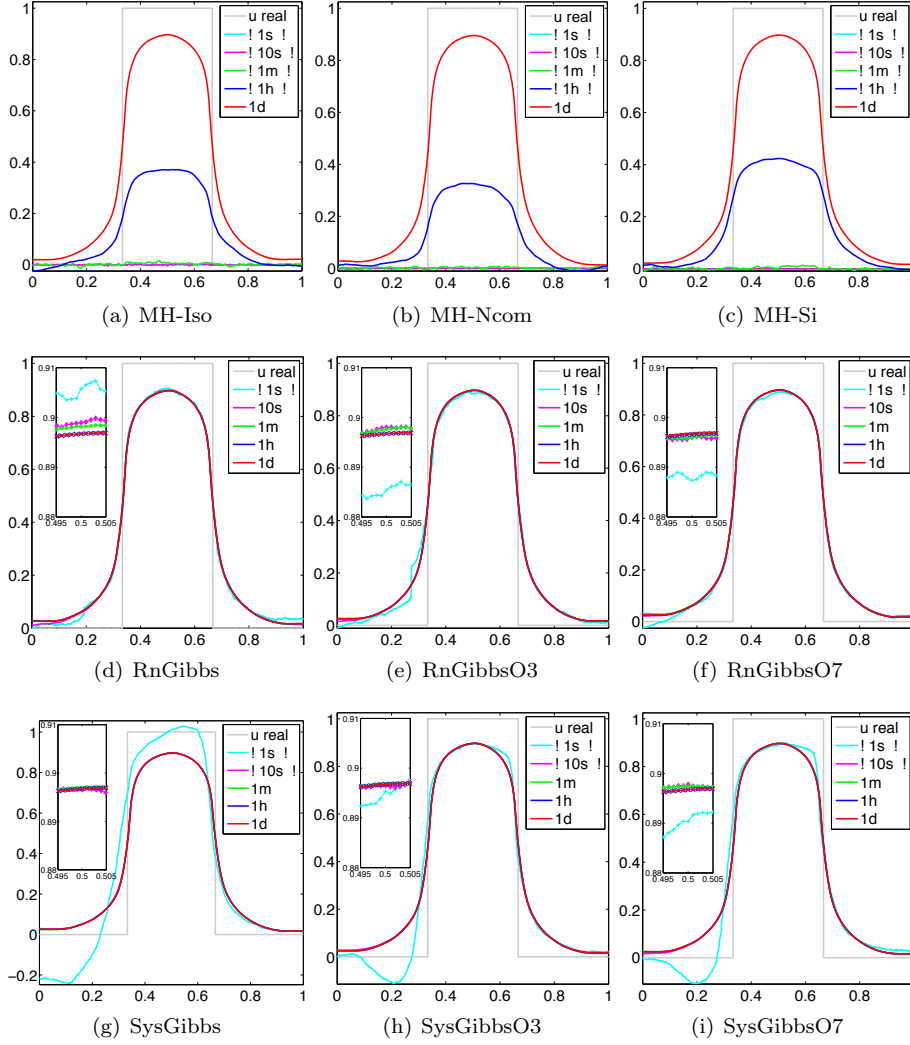


Figure 8. Visual impressions of the CM estimate obtained by a sampler after a certain computation time ($n = 1023$). Computation times, which are framed by exclamation marks, indicate that the burn-in time listed in Table 1 was not reached yet. For the Gibbs samplers, a zoom into the interval $[0.495, 0.505]$ is added.

intervals and the up- and down-scaling factors have been chosen carefully to optimize the performance of the samplers while keeping the adaptation stable, i.e., monotonic. The forward mapping is implemented using ffts. For the Gibbs sampler, we note that $V = I_n$ and that (19) simplifies to

$$(\psi_i^t \Psi_{[-i]}) \xi_{[-i]} = \frac{1}{2\sigma^2} [(A^t A) \cdot \xi]_i - \xi_i \|\psi_i\|_2^2, \quad (32)$$

which can be implemented in a efficient, direct way.




Figure 9. CM estimates for growing n and $\lambda_n \propto \sqrt{n+1}$. Sampling was performed with the RnGibbs sampler.




Figure 10. Autocorrelation plots for random and systematic scan Gibbs sampling using different values for oriented overrelaxation parameter N_O . Red \diamond : O1 (i.e., no overrelaxation at all); green \blacktriangle : O3; orange \blacktriangledown : O7; pink \blacksquare : O13; blue \bullet : O21.

3.2.2. Visual Results The practical procedure to compute visual results is identical to the one used in Section 3.1.4. In Figures 12-14, the CM estimates computed after 1, 5 and 20 hours are shown. The results of MH-Ncomp, RnGibbsO3 and SysGibbsO3 are omitted here. Choosing a good scaling to compare the results for a single sampling method is not easy because of outliers in the 1h image. These outliers would either lower the contrast if a simple linear min-max scaling based on all images is chosen or would lead to the impression that constant regions are growing if an individual scaling for each image is used. The scaling we used is generated in the following way: For each method, we merged and sorted the pixel values of all three CM estimates. From this sorted set, the smallest and largest values are discarded, using 0.1% and 99.9% as thresholds, respectively. A linear min-max scaling is generated from the remaining values, and the discarded values are mapped to the beginning and end to this color




Figure 11. Left: The colour scale used in all 2D images. Middle: The unknown function \tilde{u} . Right: The measurement data m .

scale, respectively.

An examination of $\log[p(u_i|m)]$ was again used to determine the burn-in steps K_0 . For the Gibbs samplers, the burn-in steps were between 12 and 30 and in all images shown, the burn-in phase was already completed. For the MH samplers, the plots of $\log[p(u_i|m)]$ suggested that even after 20 hours of computation, the chain was still far away from the central parts of the posterior, which is also evident from the CM estimates.

4. Discussion and Conclusions

4.1. MH-Samplers

For the specific scenario we examined, the efficiency of the basic MH samplers decreases when either the influence of the L1-type prior increases (i.e., λ is increased) or the dimension of the unknowns, n , is increased. Figures 3(a), 4(a), 5 and 7 and Tables 1 and 2 clearly document this. The largest number of unknowns examined was $n = 1\,023$, which is still moderate for typical inverse problem scenarios. However, for this number of unknowns, both the burn-in time and the time to decrease the autocorrelation of a new sample below 1% are in the order of a few hours. In Figures 8(a) - 8(c), this is visualized by the slow convergence of computed CM estimates in the 1D scenario. A (visually) satisfactory result is only obtained after 1 day of computation time. In the 2D example, no satisfactory result could be obtained, even after 20 hours of computation time (see Figure 12). The examination of $\log[p(u_i|m)]$ suggested that the computation required to obtain such a result is of orders larger. In total, our detailed studies support the empirical findings of former applications of basic MH-samplers to L1-type priors, see, e.g., [36, 33].




Figure 12. Visual results for image deblurring with an impulse prior in 2D, part 1. In the upper left corner of each subfigure, a zoom into the marked area in the original figure is shown. The scaling used in these images is explained in the text.

4.2. Gibbs-Samplers

We again stress that the Gibbs samplers we proposed and examined have to be considered as very basic variants of Gibbs sampling (cf. Section 2.2). This makes it even more surprising that for these samplers show totally different trends compared to the MH-samplers.

Random Scan Gibbs Samplers: For RnGibbs, RnGibbsO3 and RnGibbsO7, the required burn-in steps stay constant when increasing λ and clearly decrease when increasing n , cf. Table 1. Even as the computational costs of drawing a new sample increases with n , this effect keeps the computational time to draw the required number of burn-in steps almost constant. Figures 3(b), 4(b) and Table 2 show that the decay of $R(\tau)$ and $R^*(t)$ is even faster for increasing λ . From Figures 5, 6 and Table 2 we see that for increasing n , this is also true for $R(\tau)$. For $R^*(t)$, we see in Figure 7 that for large n the temporal decay cannot further decrease. This is a normal saturation effect because the autocorrelation decrease is bounded. It would even occur for an i.i.d. series of n dimensional random variables, if the computation time would increase with n . The visual results (see Figures 8(d) - 8(f), 13,14) clearly support these findings. Especially the short burn-in times are noticeable. In both scenarios, the CM estimate using the shortest computation time already represents the most




Figure 13. Visual results for image deblurring with an impulse prior in 2D, part 2. In the upper left corner of each subfigure, a zoom into the marked area in the original figure is shown. The scaling used in these images is explained in the text.

important features of the final solution. Using oriented overrelaxation in combination with random scan Gibbs sampling does not seem to lead to any problems concerning the ergodicity of the chain. The autocorrelation plots in Figures 3(b), 6 and 10(a) are still monotonic and positive. Clearly, the decay of $R(\tau)$ is faster using overrelaxation. Concerning $R^*(t)$, overrelaxation is only effective, if the additional computational cost is negligible compared to other parts of the sampling process.

Systematic Scan Gibbs Samplers: For SysGibbs, SysGibbsO3 and SysGibbsO7, the results are less clear. At first glance, the trends in their results seem to be rather similar to the random scan samplers and within a direct comparison, they often seem to outperform them, see, e.g., Table 2. However, Figures 3(b), 6 and 10(b) show that for growing n and N_O , the plots of $R(\tau)$ start to oscillate and are clearly negative in some areas. It seems that in combination with the TV prior, the subsequent update of neighboring increments leads to non-ergodic tendencies in the sampling procedure. These tendencies are amplified when using oriented overrelaxation. Producing anti-correlated samples may in fact be advantageous for certain tasks [37]. However, we would, in general, not advise to use systematic Gibbs sampling. The additional computational cost of drawing a random component to update in each step is small. In contrast, a sampler that relies on a non-ergodic mapping may produce unpredictable results for




Figure 14. Visual results for image deblurring with an impulse prior in 2D, part 3. In the upper left corner of each subfigure, a zoom into the marked area in the original figure is shown. The scaling used in these images is explained in the text.

certain tasks.

4.3. General

There are multiple reasons for the loss of performance of the basic MH samplers compared to the basic Gibbs samplers in the specific scenarios we examined. The crucial part for an MH sampler is the design of a good proposal distribution. As explained in Section 2.2, the basic MH samplers we applied are “black-box sampler” algorithms. In the design of their proposal distributions, no specific information about the posterior was taken into account. In return, they usually exhibit very fast computation times. The standard proposal distributions we used are designed to sample from low dimensional, Gaussian-like distributions. However, high dimensional posteriors from sparsity promoting priors have very different properties. Standard MH-samplers have to take very small steps to obtain a good acceptance rate. This leads to long burn-in times and a slow decrease in autocorrelation. The situation is similar with optimization algorithms used for MAP estimation (cf., Section 2.1). Black-box optimization algorithms that only rely on evaluating the objective function (i.e., $\log[p(u|m)]$) are usually too slow when applied to specific, high dimensional posteriors. The basic Gibbs samplers we proposed incorporate more posterior-specific information into the sampling procedure at the costs of a larger computation time. The

conditional single component densities, which can be regarded as optimal transition kernels, are computed and sampled from explicitly. This small extra amount of incorporating problem specific information already seems to be sufficient to generate very promising sampling procedures for high dimensional Bayesian inversion using L1-type priors (the dimensions of the unknowns used in Figure 9 and the 2D scenario are far beyond any previously reported use of MCMC for L1-type inverse problems). In general, both sampling techniques have advantages and disadvantages and will outperform the other given a specific scenario.

For the very reason that we only used very basic Gibbs samplers, our results also challenge common beliefs about the feasibility of MCMC sampling in high dimensional inverse problems in general. We showed that MCMC schemes are not in general slow and scale bad with increasing dimension. We rather think that MCMC schemes for inverse problems are far less elaborate compared to optimization schemes up to now. With regard to the corresponding optimization algorithms for MAP estimates, one possible reason for the superior performance of the single component Gibbs samplers might be the transformation of the posterior into the basis v_1, \dots, v_n (cf., 2.3). In this basis, the prior diagonalizes. It can be shown that the MAP estimate is sparse in this basis, i.e., many basis coefficients are exactly zero. Many optimization algorithms to compute the MAP estimate take advantage of this and perform better when transformed into that basis. One could argue that this might be the case for the sampling procedures as well, and that a fair comparison between MH and Gibbs samplers would need to transform the MH samplers into the basis v_1, \dots, v_n as well. However, there are reasons why this argument is not valid. The striking advantage of MH samplers is their simple, “black-box”-like implementation. In practice, they are normally implemented in the most direct way and we stuck to that paradigm. In addition, in the 2D case, $v_i = e_i$, i.e., both MH and Gibbs samplers are already formulated in the right basis. However, this does not affect the bad performance of the MH samplers compared to the Gibbs samplers. But most importantly, while the MAP estimate is sparse, the CM estimate is not, and single samples from the posterior are not sparse as well. Theoretically, it has been shown in [38] that for denoising using a TV prior, the CM estimate is, in fact, never sparse. One can see this, e.g., in Figure 8. The estimates are neither sparse in the normal basis, nor in the increment basis. In Figures 12- 14, one can clearly see that this is similar for the normal L1 prior.

4.4. Outlook and Extensions

In this first study, we only compared very basic variants of MH and Gibbs sampling. In the future, a comparison to more sophisticated variants of MH schemes such as *delayed rejection* [40, 23], *adaptive Metropolis* schemes [24, 23] or the *t-walk* [9] has to be undertaken. The most promising technique for our scenario might be to combine a tailored variant of delayed rejection with SCAM [24]. Many of the sophisticated MH variants have been developed for the study of non-linear, computationally extensive and high dimensional inverse problems (see [13] for a recent overview), i.e., situations where MH sampling is the only MCMC technique that can be applied. Improving the basic Gibbs sampling schemes used here by adding adaptive elements or optimal directions is far less developed until now [10] and is an interesting future topic of research. A comparison between sophisticated (possibly adaptive, i.e., non-markovian) variants of MH and Gibbs sampling will also need a concrete application scenario since a general comparison by the measures used in this article is less meaningful.

We assumed that the variance σ^2 of the noise term (cf. Section 2.1) is known exactly (or a good estimate is available), which is not always the case in practical applications. The Bayesian framework can account for the uncertainty of this model parameter as well: σ^2 is treated like the other unknowns u (but assumed to be independent from them) and the available information about its typical values are expressed by a prior $p_{pr}(\sigma^2)$. An advantageous choice for $p_{pr}(\sigma^2)$ is given by the *conjugate prior* w.r.t. $p_{li}(m|u, \sigma^2)$, which is the *inverse gamma distribution* [19]:

$$p(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} \exp\left(-\frac{\beta}{x}\right), \quad (33)$$

with *shape* and *scale* parameters α and β (Γ denotes the gamma function). Now the joint posterior for u and σ^2 given the data m reads:

$$\begin{aligned} p_{post}(u, \sigma^2|m) &\propto \\ &\left(\frac{1}{2\sigma^2}\right)^{\frac{k}{2}} \exp\left(-\frac{1}{2\sigma^2}\|m - Au\|_2^2 - \lambda|Du|\right) \exp\left(-(\alpha+1)\log(\sigma^2) - \frac{\beta}{\sigma^2}\right) \\ &= \exp\left(-\frac{\frac{1}{2}\|m - Au\|_2^2 + \beta}{\sigma^2} - \lambda|Du| - (\alpha+1+k/2)\log(\sigma^2)\right) \end{aligned} \quad (34)$$

A comparison with (33) shows that the conjugacy property of the prior has the effect that the posterior of σ^2 conditioned on both m and u is, again, an inverse gamma distribution with shape and scale parameters $\tilde{\alpha}$ and $\tilde{\beta}$ given by :

$$\tilde{\alpha} = \alpha + k/2; \quad \tilde{\beta} = \frac{1}{2}\|m - Au\|_2^2 + \beta \quad (35)$$

This allows us to perform Bayesian inference for the joint posterior using Gibbs sampling: For sampling along a component of u we can use the fast samplers presented here§, and for sampling over σ^2 conditioned on all other parameters we can use standard implementations of gamma samplers. Such a Gibbs sampler can, e.g., be used to infer a joint CM estimate $(u_{jCM}, \sigma_{jCM}^2)$. The information given by σ_{jCM}^2 can be used to evaluate or improve the measurement setup or to inform other Bayesian reconstructions. The estimate u_{jCM} compared to an estimate assuming a single, constant σ^2 contains the *marginalized* uncertainty about σ^2 and may yield a more robust estimate of u in practical applications. In principle, it is possible to marginalize over σ^2 explicitly (the computation is similar to [19]):

$$p_{post}(u|m) = \int p_{post}(u, \sigma^2|m) d\sigma^2 \propto \left(\frac{1}{2}\|m - Au\|_2^2 + \beta\right)^{-(\alpha+k/2)} p_{pr}(u) \quad (36)$$

$$\propto \left(1 + \frac{t(u)^2}{\nu}\right)^{-\frac{1}{2}(\nu+1)} p_{pr}(u), \quad (37)$$

$$\text{where } \nu = 2\alpha + 1 + k; \quad t(u) = \left(\frac{\frac{1}{2}\|m - Au\|_2^2}{\beta(2\alpha + 1 + k)^{1/2}}\right)^{1/2}. \quad (38)$$

This is a (one sided) Student's t-distribution for $t(u)$ with ν *degrees of freedom*. However, working with this distribution directly is more difficult since it is not *log-concave* (this problem gets worse if an individual variance σ_i^2 for each noise channel is assumed). Working with the full joint posterior instead can circumvent some of these problems.

§ The practical implementation has to be slightly adopted in the sense that the varying σ^2 has to be removed from all expressions that are precomputed and added back to them at runtime.

The Gibbs samplers presented here have to be generalized to work with arbitrary D , e.g., to deal with anisotropic total variation priors in arbitrary dimensions and with arbitrary boundary conditions. In addition, an extension to *block sparse* priors, which rely on mixed L2-L1 norms, would be advantageous to, e.g., address isotropic total variation priors.

Parallelization of MCMC sampling is easily implemented. In the most basic form, N independent chains are generated, each on one CPU. However, the efficiency of this approach is strongly limited by the burn-in and mixing time [37]. If all chains are initialized at the same state, parallelization is only efficient, if the chains become independent very fast. Our results suggest that parallelization of Gibbs samplers will be way more efficient than of MH samplers.

The Gibbs sampling algorithms developed by us are fast enough to tackle sampling for Bayesian inversion techniques in real applications, which will be an important topic of future work. In many applications like, e.g., limited angle CT, exploring the full range of Bayesian inversion by also incorporating sample based analysis was, up to now, rather regarded as a theoretical option, see, e.g., [49, 34].

In addition, theoretical questions concerning sparse Bayesian inversion, like, e.g., the ones addressed in [36, 35] can be also be addressed numerically (cf. Section 3.1.4).

Acknowledgments

We would like to thank the anonymous referees and Martin Burger for their helpful critiques and comments that significantly improved this manuscript.

Appendix A. Code

On the authors homepage^{||} Matlab code supporting this publication is provided. It contains scripts to create the scenarios examined in the numerical studies as well as implementations of all the Gibbs sampling algorithms presented here. One should, however, mention that Matlab is not very well suited for these implementation as the sampling algorithms consist of very sequential but rather basic procedures. We therefore also provide alternative implementations of the samplers using Fortran within `.mex`-files, which can be compared to the corresponding `.m`-files. The speed-up factor for the RnGibbs Sampler ranges from 43 to 18 using $n = 63$ or $n = 1023$ for the 1D scenario and is 1.35 for the 2D scenario with $n = 261\,121$ (the computation of b is by far the most expensive computational task, and little gain can be expected from a direct implementation compared to Matlab).

Appendix B. Implementation

In this section, we give details on how to implement formulas (21), (22) and (23). The complementary error function and its inverse are difficult to handle numerically because there are no identities that allow to rescale or shift their evaluation to other intervals. For the applications we address, problems due to limited precision occur when formulas (21), (22) and (23) are implemented directly (formula (21) is only required for applying ordered overrelaxation). Dependent on the signs of α_+ and α_- , we use different alternative formulas that allow for a stable numerical evaluation.

^{||} Currently: <http://wwwmath.uni-muenster.de/num/burger/organization/lucka>

Additionally, we express $\operatorname{erfc}(x)$ in terms of the scaled complementary error function $\operatorname{erfcx}(x) = \exp(x^2)\operatorname{erfc}(x)$, which decays less fast for $x \rightarrow +\infty$. We only list the results here (the corresponding transformations are elementary but lengthy to write down). Because $c \geq 0$, not both α_+ and α_- can be negative, which leaves three different cases to examine:

$\alpha_+ > 0, \alpha_- > 0$: Let $\gamma_{++} := [\operatorname{erfcx}(\alpha_+) + \operatorname{erfcx}(\alpha_-)]$. Then, the parts of (21) are given by:

$$y < 0 : \exp(-ay^2 + 2\sqrt{ay}\alpha_+) \operatorname{erfcx}(-\sqrt{ay} + \alpha_+)/\gamma_{++} \quad (\text{B.1})$$

$$y > 0 : 1 - \exp(-ay^2 - 2\sqrt{ay}\alpha_-) \operatorname{erfcx}(\sqrt{ay} + \alpha_-)/\gamma_{++} \quad (\text{B.2})$$

The arguments of $\operatorname{erfcinv}$ in (22) and (23) are given by:

$$\text{In (22) : } r \exp(-\alpha_+^2) \gamma_{++} \quad (\text{B.3})$$

$$\text{In (23) : } (1-r) \exp(-\alpha_-^2) \gamma_{++} \quad (\text{B.4})$$

$\alpha_+ < 0, \alpha_- > 0$: Since erfcx increases very fast for $x \rightarrow -\infty$ one has to use the identity $\operatorname{erfcx}(-x) = 2 \exp(x^2) - \operatorname{erfcx}(x)$. Let $\gamma_{-+} := [\operatorname{erfcx}(-\alpha_+) - \operatorname{erfcx}(\alpha_-)]$. Then, the formulas to implement the parts of (21) are given by:

$$\begin{aligned} y < 0 : & \frac{\exp[-(\sqrt{ay} - \alpha_+)^2] \operatorname{erfcx}(-\sqrt{ay} + \alpha_+)}{2 - \exp(-\alpha_+^2) \gamma_{-+}} \quad (\text{B.5}) \\ -\sqrt{ay} + \alpha_+ > 0 : & \end{aligned}$$

$$\begin{aligned} y < 0 : & \frac{\{2 - \exp[-(\sqrt{ay} - \alpha_+)^2] \operatorname{erfcx}(\sqrt{ay} - \alpha_+)\}}{2 - \exp(-\alpha_+^2) \gamma_{-+}} \quad (\text{B.6}) \\ -\sqrt{ay} + \alpha_+ < 0 : & \end{aligned}$$

$$\begin{aligned} y > 0 : & 1 - \frac{\exp[-(\sqrt{ay} - \alpha_-)^2] \operatorname{erfcx}(\sqrt{ay} + \alpha_-)}{2 \exp(\frac{bc}{a}) - \exp(-\alpha_-^2) \gamma_{-+}} \quad (\text{B.7}) \\ \sqrt{ay} + \alpha_- > 0 : & \end{aligned}$$

$$\begin{aligned} y > 0 : & 1 - \frac{2 - \exp[-(\sqrt{ay} - \alpha_-)^2] \operatorname{erfcx}(-\sqrt{ay} - \alpha_-)}{2 \exp(\frac{bc}{a}) - \exp(-\alpha_-^2) \gamma_{-+}} \quad (\text{B.8}) \\ \sqrt{ay} + \alpha_- < 0 : & \end{aligned}$$

The arguments of $\operatorname{erfcinv}$ in (22) and (23) are given by:

$$\text{In (22) : } r [2 - \exp(-\alpha_+^2) \gamma_{-+}] \quad (\text{B.9})$$

$$\text{In (23) : } (1-r) \left[2 \exp\left(\frac{bc}{a}\right) - \exp(-\alpha_-^2) \gamma_{-+} \right] \quad (\text{B.10})$$

$\alpha_+ > 0, \alpha_- < 0$: Let $\gamma_{+-} := [\operatorname{erfcx}(\alpha_+) - \operatorname{erfcx}(-\alpha_-)]$. Then, the parts of (21) are given by:

$$y < 0 : \frac{\exp[-(\sqrt{ay} - \alpha_+)^2] \operatorname{erfcx}(-\sqrt{ay} + \alpha_+)}{2 \exp(-\frac{bc}{a}) + \exp(-\alpha_+^2) \gamma_{+-}} \quad (\text{B.11})$$

$$\begin{aligned} y > 0 : & 1 - \frac{\exp[-(\sqrt{ay} + \alpha_-)^2] \operatorname{erfcx}(\sqrt{ay} + \alpha_-)}{2 + \exp(-\alpha_-^2) \gamma_{+-}} \quad (\text{B.12}) \\ \sqrt{ay} + \alpha_- > 0 : & \end{aligned}$$

$$\begin{aligned} y > 0 : & 1 - \frac{2 - \exp[-(\sqrt{ay} + \alpha_-)^2] \operatorname{erfcx}(-\sqrt{ay} - \alpha_-)}{2 + \exp(-\alpha_-^2) \gamma_{+-}} \quad (\text{B.13}) \\ \sqrt{ay} + \alpha_- < 0 : & \end{aligned}$$

$$(\text{B.14})$$

The arguments of erfcinv in (22) and (23) are given by:

$$\text{In (22)} : \quad r \left[2 \exp \left(\frac{-bc}{a} \right) + \exp(-\alpha_+^2) \gamma_{+-} \right] \quad (\text{B.15})$$

$$\text{In (23)} : \quad (1-r) [2 + \exp(-\alpha_-^2) \gamma_{+-}] \quad (\text{B.16})$$

Using the above expressions directly can still lead to stability issues, because very large numbers are often multiplied with very small numbers. It is preferable to compute the logarithms of the expressions first. For this, let $x > 0$, $(x+y) > 0$ then:

$$\log(x+y) = \log(x) + \log(1 + \text{sign}(y) \exp(\log(|y|) - \log(x))) \quad (\text{B.17})$$

Using this identity we can compute the logarithms of expressions (B.1) - (B.16). We note that $\text{sign}[\pm \text{erfcx}(\cdot)] = \pm 1$.

$$\log[(B.1)] = (-ay^2 + 2\sqrt{ay}\alpha_+) + \log[\text{erfcx}(-\sqrt{ay} + \alpha_+)] - \log(\gamma_{++}) \quad (\text{B.18})$$

$$\log[1 - (B.2)] = (-ay^2 - 2\sqrt{ay}\alpha_-) + \log[\text{erfcx}(\sqrt{ay} + \alpha_-)] - \log(\gamma_{++}) \quad (\text{B.19})$$

$$\log[(B.3)] = \log(r) - \alpha_+^2 + \log(\gamma_{++}) \quad (\text{B.20})$$

$$\log[(B.4)] = \log(1-r) - \alpha_-^2 + \log(\gamma_{++}) \quad (\text{B.21})$$

$$\begin{aligned} \log[(B.5)] = & -(\sqrt{ay} + \alpha_+)^2 + \log[\text{erfcx}(-\sqrt{ay} + \alpha_+)] - \log(2) \\ & - \log\{1 - \text{sign}(\gamma_{+-}) \exp[-\alpha_+^2 + \log(|\gamma_{+-}|) - \log(2)]\} \end{aligned} \quad (\text{B.22})$$

$$\begin{aligned} \log[(B.5)] = & \log\left(1 - \exp\left\{\log[\text{erfcx}(\sqrt{ay} - \alpha_+)] - \log(2) - (\sqrt{ay} - \alpha_+)^2\right\}\right) \\ & - \log\{1 - \text{sign}(\gamma_{+-}) \exp[-\alpha_+^2 + \log(|\gamma_{+-}|) - \log(2)]\} \end{aligned} \quad (\text{B.23})$$

$$\begin{aligned} \log[1 - (B.7)] = & -(\sqrt{ay} + 2\alpha_-)^2 + \log[\text{erfcx}(\sqrt{ay} + \alpha_-)] - \log(2) - \frac{bc}{a} \\ & - \log\left\{1 - \text{sign}(\gamma_{+-}) \exp\left[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2) - \frac{bc}{a}\right]\right\} \end{aligned} \quad (\text{B.24})$$

$$\begin{aligned} \log[1 - (B.8)] = & \log\left(1 - \exp\left\{\log[\text{erfcx}(-\sqrt{ay} - \alpha_-)] - \log(2) - (\sqrt{ay} + \alpha_-)^2\right\}\right) \\ & - \frac{bc}{a} - \log\left\{1 - \text{sign}(\gamma_{+-}) \exp\left[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2) - \frac{bc}{a}\right]\right\} \end{aligned} \quad (\text{B.25})$$

$$\begin{aligned} \log[(B.9)] = & \log(r) + \log(2) \\ & + \log\{1 - \text{sign}(\gamma_{+-}) \exp[-\alpha_+^2 + \log(|\gamma_{+-}|) - \log(2)]\} \end{aligned} \quad (\text{B.26})$$

$$\begin{aligned} \log[(B.10)] = & \log(1-r) + \log(2) + \frac{bc}{a} \\ & + \log\left\{1 - \text{sign}(\gamma_{+-}) \exp\left[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2) - \frac{bc}{a}\right]\right\} \end{aligned} \quad (\text{B.27})$$

$$\begin{aligned} \log[(B.11)] = & -(-\sqrt{ay} + \alpha_+)^2 + \log[\text{erfcx}(-\sqrt{ay} + \alpha_+)] - \log(2) + \frac{bc}{a} \\ & - \log\left\{1 + \text{sign}(\gamma_{+-}) \exp\left[-\alpha_+^2 + \log(|\gamma_{+-}|) - \log(2) + \frac{bc}{a}\right]\right\} \end{aligned} \quad (\text{B.28})$$

$$\begin{aligned} \log[(B.12)] = & -(\sqrt{ay} + \alpha_-)^2 + \log[\text{erfcx}(\sqrt{ay} + \alpha_-)] - \log(2) \\ & - \log\{1 + \text{sign}(\gamma_{+-}) \exp[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2)]\} \end{aligned} \quad (\text{B.29})$$

$$\begin{aligned} \log[(B.13)] = & \log\left(1 - \exp\left\{\log[\text{erfcx}(-\sqrt{ay} - \alpha_-)] - \log(2) - (\sqrt{ay} + \alpha_-)^2\right\}\right) \\ & - \log\{1 + \text{sign}(\gamma_{+-}) \exp[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2)]\} \end{aligned} \quad (\text{B.30})$$

$$\begin{aligned} \log[(B.15)] &= \log(r) + \log(2) - \frac{bc}{a} \\ &\quad + \log \left\{ 1 + \text{sign}(\gamma_{+-}) \exp \left[-\alpha_+^2 + \log(|\gamma_{+-}|) - \log(2) + \frac{bc}{a} \right] \right\} \end{aligned} \quad (B.31)$$

$$\begin{aligned} \log[(B.16)] &= \log(1-r) + \log(2) \\ &\quad + \log \left\{ 1 + \text{sign}(\gamma_{+-}) \exp \left[-\alpha_-^2 + \log(|\gamma_{+-}|) - \log(2) \right] \right\} \end{aligned} \quad (B.32)$$

Now, for (22) and (23), if w denotes the logarithm of the argument of erfcinv , one can compute $\text{erfcinvlog}(w) := \text{erfcinv}[\exp(w)]$ using a standard implementation of erfcinv if w is not too small (the loss of precision using $\exp(w)$ instead of computing the full argument of erfcinv is negligible since the variation of erfcinv is very small even on logarithmic scale). However, even using 64 bit precision is not sufficient for the applications we address. Therefore, we use an asymptotic approximation of $\text{erfcinvlog}(w)$ for $w < -680$ from [1].

An approximation of $z = \text{erfcinv}[\exp(w)]$ for $w \rightarrow -\infty$ is given by:

$$\begin{aligned} \theta &:= -\log(\pi) - \log(-w) \\ v &:= (-\theta - 2) \\ s &:= 2/(\theta - 2w) \\ a_2 &:= \frac{1}{8}v \\ a_3 &:= -\frac{1}{32}(v^2 + 6v - 6) \\ a_4 &:= \frac{1}{384}(4v^3 + 27v^2 + 108v - 300) \\ z &\approx s^{-1/2} + a_2s^{3/2} + a_3s^{5/2} + a_4s^{7/2} \end{aligned} \quad (B.33)$$

The discrepancy of this approximation to the implementation of erfcinv in Matlab is $2.34 \cdot 10^{-12}$ for $w = -690$ and as it is an asymptotic formula, the error further decreases for $w \rightarrow -\infty$.

References

- [1] Digital Library of Mathematical Functions. National Institute of Standards and Technology from <http://dlmf.nist.gov/>, 2011.
- [2] J. Bardsley, D. Calvetti, and E. Somersalo. Hierarchical regularization for edge-preserving reconstruction of PET images. *Inverse Probl*, 26:035010 (16pp), 2010.
- [3] N. Bissantz, T. Hohage, A. Munk, and F. Ruymgaart. Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM J Numer Anal*, 45(6):2610–2636, 2007.
- [4] S. Brooks and G. Roberts. Assessing convergence of Markov chain Monte Carlo algorithms. *Stat Comput*, 8(4):319–335, 1998.
- [5] M. Burger, M. Moeller, M. Benning, and S. Osher. An Adaptive Inverse Scale Space Method for Compressed Sensing. CAM Report 11-08, UCLA, 2011.
- [6] M. Burger and S. Osher. A guide to the TV zoo. 2012.
- [7] D. Calvetti and E. Somersalo. A Gaussian hypermodel to recover blocky objects. *Inverse Probl*, 23(2):733–754, 2007.
- [8] D. Calvetti and E. Somersalo. Hypermodels in the Bayesian imaging framework. *Inverse Probl*, 24(3):034013 (20pp), 2008.
- [9] J. A. Christen and C. Fox. A General Purpose Sampling Algorithm for Continuous Distributions (the t-walk). *Bayesian Anal*, 5(2):263–282, 2010.
- [10] J. A. Christen, C. Fox, D. A. Pérez-Ruiz, and M. Santana-Cibrian. On optimal direction gibbs sampling. (arXiv:1205.4062), 2012.

- [11] S. Comelli. A Novel Class of Priors for Edge-Preserving Methods in Bayesian Inversion. Master's thesis, University of Milan, Italy, 2011.
- [12] M. Cowles and B. Carlin. Markov chain Monte Carlo convergence diagnostics: a comparative review. *J Am Stat Assoc*, pages 883–904, 1996.
- [13] T. Cui, C. Fox, and M. J. O'Sullivan. Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm. *Water Resour Res*, 47, 2011.
- [14] D. L. Donoho. Compressed sensing. *IEEE Trans Inf Theory*, 52(4):1289–1306, 2006.
- [15] Y. C. Eldar and G. Kutyniok, editors. *Compressed Sensing - Theory and Applications*. Cambridge University Press, New York, 1st edition, May 2012.
- [16] T. Eltoft, T. Kim, and T.-W. Lee. On the multivariate Laplace distribution. *IEEE Signal Process Lett*, 13(5):300–303, may 2006.
- [17] H. Engl, M. Hanke-Bourgeois, and A. Neubauer. *Regularization of Inverse Problems. Mathematics and Its Applications*. Springer Netherland, Berlin, 1996.
- [18] A. E. Gelfand and A. F. M. Smith. Sampling-based approaches to calculating marginal densities. *J Am Stat Assoc*, 85(410):398–409, 1990.
- [19] A. Gelman. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, 1(3):515–533, 2006.
- [20] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Pattern Anal Mach Intell*, PAMI-6(6):721–741, nov. 1984.
- [21] T. Goldstein and S. Osher. The Split Bregman method for L1-regularized problems. *SIAM J Img Sci*, 2:323–343, April 2009.
- [22] H. Haario, M. Laine, M. Lehtinen, E. Saksman, and J. Tamminen. Markov Chain Monte Carlo Methods for High Dimensional Inversion in Remote Sensing. *J R Stat Soc Series B Stat Methodol*, 66(3):591–607, 2004.
- [23] H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM: Efficient adaptive MCMC. *Stat Comput*, 16(4):339–354, December 2006.
- [24] H. Haario, E. Saksman, and J. Tamminen. Componentwise adaptation for high dimensional MCMC. *Comput Stat*, 20(2):265–273, June 2005.
- [25] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [26] T. Helin. On infinite-dimensional hierarchical probability models in statistical inverse problems. *Inverse Probl Imaging*, 3:567–597, 2010.
- [27] T. Helin and M. Lassas. Hierarchical Models in Statistical Inverse Problems and the Mumford–Shah Functional. Technical Report arXiv:0908.3396, Aug 2009.
- [28] D. Higdon, H. Lee, and C. Holloman. Markov Chain Monte Carlo-based Approaches for Inference in Computationally Intensive Inverse Problems. 2003.
- [29] A. Hofinger. *Ill-posed problems: Extending the deterministic theory to a stochastic setup*. PhD thesis, Universität Linz, 2006.
- [30] J. Idier. *Bayesian Approach to Inverse Problems*. Wiley-ISTE, 2008.
- [31] J. P. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*, volume 160 of *Applied Mathematical Sciences*. Springer New York, 2005.
- [32] A. Klenke. *Probability Theory: A Comprehensive Course*. Springer London, 1st edition, 2008.
- [33] V. Kolehmainen, M. Lassas, K. Niinimäki, and S. Siltanen. Sparsity-promoting Bayesian inversion. *Inverse Probl*, 28(2):025005 (28pp), 2012.
- [34] V. Kolehmainen, S. Siltanen, S. Järvenpää, J. P. Kaipio, P. Koistinen, M. Lassas, J. Pirttilä, and E. Somersalo. Statistical inversion for medical x-ray tomography with few radiographs: II. Application to dental radiology. *Phys Med Biol*, 48(10):1465–90, May 2003.
- [35] M. Lassas, E. Saksman, and S. Siltanen. Discretization invariant Bayesian inversion and Besov space priors. *Inverse Probl Imaging*, (3):87–122, Jan 2009.
- [36] M. Lassas and S. Siltanen. Can one use total variation prior for edge-preserving Bayesian inversion? *Inverse Probl*, 20:1537–1563, 2004.
- [37] J. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer New York, 2008.
- [38] C. Louchet. *Variational and Bayesian models for image denoising: from total variation towards non-local means*. PhD thesis, Université Paris Descartes, 2008.
- [39] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J Chem Phys*, 21:1087–1092, 1953.
- [40] A. Mira. On Metropolis-Hastings algorithms with delayed rejection. *Metron*, LIX(3-4):231–241, 2001.
- [41] J. D. Moulton, C. Fox, and D. Svyatskiy. Multilevel approximations in sample-based inversion

- from the Dirichlet-to-Neumann map. *J Phys Conf Ser*, 124(1):012035, 2008.
- [42] P. Neal and G. Roberts. Optimal scaling for partially updating MCMC algorithms. *Ann Appl Probab*, 16(2):475–515, 2006.
- [43] R. M. Neal. Suppressing Random Walks in Markov Chain Monte Carlo Using Ordered Overrelaxation. Technical report, Learning in Graphical Models, 1995.
- [44] A. Parker and C. Fox. Sampling Gaussian distributions in Krylov spaces with conjugate gradients. *SIAM J Sci Comput*, 2012.
- [45] G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Stat Sci*, 16(4):351–367, 2001.
- [46] G. O. Roberts and S. K. Sahu. Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler. *Journal of the Royal Statistical Society.*, 59(2):pp. 291–317, 1997.
- [47] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D Nonlinear Phenomena*, 60:259–268, November 1992.
- [48] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003.
- [49] S. Siltanen, V. Kolehmainen, S. Järvenpää, J. P. Kaipio, P. Koistinen, M. Lassas, J. Pirttilä, and E. Somersalo. Statistical inversion for medical x-ray tomography with few radiographs: I. General theory. *Phys Med Biol*, 48(10):1437–63, May 2003.
- [50] A. Tarantola and B. Valette. Inverse Problems = Quest for Information. *J Geophys*, 50:159–170, 1982.
- [51] M. Thompson. A comparison of methods for computing autocorrelation time. *Arxiv preprint arXiv:1011.0175*, 2010.
- [52] U. Toussaint. Bayesian inference in physics. *Rev Mod Phys*, 83(3):943–999, 2011.